

What is claimed is:

1 Claim 1. A multiprocessor computer system comprising,
2 a computer system having a plurality of processing
3 nodes and caches and a node controller which use processor
4 state information according to mappings provided by
5 supervisor software or firmware of allowable physical
6 processors to an application workload to determine which
7 coherent cache regions in the system are required to examine
8 a coherency transaction produced by a single originating
9 processor's storage request.

1 Claim 2. The multiprocessor computer system according to
2 claim 1 wherein a node of the computer has dynamic coherency
3 boundaries such that the hardware uses only a subset of the
4 total processors in the system for a single workload at any
5 specific point in time and optimizes the cache coherency as
6 the supervisor software or firmware expands and contracts
7 the number of processors which are being used to run any
8 single workload.

1 Claim 3. The multiprocessor computer system according to
2 claim 1 wherein multiple instances of a node are connected
3 with a second level controller to create a large
4 multiprocessor system.

1 Claim 4. The multiprocessor computer system according to
2 claim 1 wherein said node controller uses mode bits to
3 determine which processors must receive any given
4 transaction that is received by the node controller.

1 Claim 5. The multiprocessor computer system according to
2 claim 1 wherein a second level controller is provided which
3 uses the mode bits to determine which nodes must receive any

4 given transaction that is received by the second level
5 controller.

1 Claim 6. The multiprocessor computer system according to
2 claim 1 wherein logical partitions are provided and
3 mapping of logical partitions to allowable physical
4 processors is provided by provided by supervisor software or
5 firmware of allowable physical processors to an application
6 workload.

1 Claim 7. The multiprocessor computer system according to
2 claim 1 wherein logical partitions are provided for the
3 supervisor software or firmware which maps allowable
4 physical processors to an application workload and a
5 hypervisor assigns cache coherence regions which encompass
6 subsets of the total number of processors and caches in the
7 system chosen for their physical proximity and defines a
8 distinct cache coherency region for each partition.

1 Claim 8. The multiprocessor computer system according to
2 claim 1 wherein a single workload uses only a subset of the
3 total processors in the computer system for a single
4 workload at any specific point in time for an assigned
5 partition and a distinct cache coherency is optimized for
6 the address space of the assigned partition as the
7 supervisor software or firmware expands and contracts the
8 number of processors which are being used to run any single
9 workload in said assigned partition.

1 Claim 9. The multiprocessor computer system according to
2 claim 1 wherein a single workload uses only a subset of the
3 total processors in the computer system for a single
4 workload at any specific point in time, and

5 multiple cache coherent regions are assigned for different
6 partitions as more independent workloads coexist on the same
7 hardware.

1 Claim 10. The multiprocessor computer system according to
2 claim 1 wherein cache coherence regions encompass subsets of
3 the total number of processors and caches in the computer
4 system and a single workload uses only a subset of the total
5 processors in the computer system for a single workload at
6 any specific point in time for an assigned partition and a
7 distinct cache coherency is optimized for the address space
8 of the assigned partition as the supervisor software or
9 firmware expands and contracts the number of processors
10 which are being used to run any single workload in said
11 assigned partition.

1 Claim 11. The multiprocessor computer system according to
2 claim 1 wherein software and/or firmware define which subset
3 of processors in a large multiprocessor must participate in
4 a coherency transaction independent of which processing node
5 is connected to the physical DRAM storage being requested by
6 said single originating processor.

1 Claim 12. The multiprocessor computer system according to
2 claim 11 wherein the movement of a process between nodes of
3 a large multiprocessor is effectuated without moving
4 physical storage contents and without requiring subsequent
5 broadcasting of the storage references originated by the
6 process from said single originating storage request to all
7 of the caches in the multiprocessor.

1 Claim 13. The multiprocessor computer system according to
2 claim 1 wherein cache coherence mode bits are appended to a

3 processors storage transactions when transmitted to a
4 connected processor of said multiprocessor computer system.

1 Claim 14. The multiprocessor computer system according to
2 claim 13 wherein said cache coherence mode bits are used in
3 a decision determining whether the single originating
4 processor's storage request must be transmitted to
5 additional processors in the system.

1 Claim 15. The multiprocessor computer system according to
2 claim 14 wherein an increase in the effective utilization of
3 the address bandwidth of the buses used to interconnect the
4 processors of a multiprocessor system allows movement of
5 workload among physical processors in a multiprocessor
6 system at the same time as a reduction of the address
7 bandwidth required to maintain cache coherency among all the
8 processors.